Noname manuscript No.
(will be inserted by the editor)

Active Target Tracking with Self-Triggered Communications in Multi-Robot Teams

Lifeng Zhou · Pratap Tokekar

arXiv:1704.07475v1 [cs.RO] 24 Apr 2017

Received: date / Accepted: date

Abstract We study the problem of reducing the amount of communication in decentralized target tracking. We focus on the scenario where a team of robots are allowed to move on the boundary of the environment. Their goal is to seek a formation so as to best track a target moving in the interior of the environment. The robots are capable of measuring distances to the target. Decentralized control strategies have been proposed in the past that guarantee that the robots asymptotically converge to the optimal formation. However, existing methods require that the robots exchange information with their neighbors at all time steps. Instead, we focus on decentralized strategies to reduce the amount of communication among robots.

We propose a self-triggered communication strategy that decides when a particular robot should seek up-to-date information from its neighbors and when it is safe to operate with possibly outdated information. We prove that this strategy converges asymptotically to a desired formation when the target is stationary. For the case of a mobile target, we propose an extension whereby each robot decides its optimal partner to share its measurements with using observability as a criterion. We evaluate all the approaches (constant communication and self-triggered communication with centralized and decentralized sensor fusion) through simulations.

Keywords multi-robot systems · target tracking · networked control

L. Zhou
Department of Electrical & Computer Engineering
Virginia Tech, U.S.A.
E-mail: lfzhou@vt.edu

P. Tokekar
Department of Electrical & Computer Engineering
Virginia Tech, U.S.A.
E-mail: tokekar@vt.edu

1 Introduction

The target tracking problem has been one of the most well-studied problems in the robotics community [1] and finds many applications such as surveillance [3, 7, 17], crowd monitoring [2, 19], and wildlife monitoring [4, 18]. In this paper, we study *active* target tracking with a team of robots where the focus is on actively controlling the state of the robot. The robots can exchange information with each other and then decide how to move, so as to best track the target. It is typically assumed that exchanging information is beneficial. It is typical to design strategies by assuming that the robots will exchange their information at each time step irrespective of whether that information is worth exchanging. In this paper, we investigate the problem of deciding when is it worthwhile for the robots to exchange information and when is it okay to use possibly outdated information.

The motivation for our work stems from the observation that communication can be costly. For example, for smaller robots radio communication can be a significant source of power consumption. The robots can extend their lifetime by reducing the time spent communicating (equivalently, number of messages sent). Our goal is thus to determine a strategy that communicates only when required without significantly affecting the tracking performance.

We study this problem in a simple target tracking scenario first introduced by Martinez and Bullo [13]. Here, the robots are restricted to move on the boundary of a convex environment. They can obtain distance measurements towards a target moving in the interior. The goal of the robots is to position themselves so as to maximize the information gained. The authors proposed a decentralized strategy where the robots com-

municate at all time steps with their neighbors and proved that it converges to the optimal (uniform) configuration. Instead, we apply a self-triggered coordination algorithm (following recent works [9, 15]) where each robot decides when to trigger communications with its neighbors. We apply this strategy to the aforementioned target tracking scenario and compare its performance relative to the constant strategy in simulations.

We also evaluate the error in tracking the target using the self-triggered and constant communication strategies. In a decentralized setting, robots can exchange information only with their neighbors. As a result, their local estimates of the target's position may differ significantly, resulting in poor tracking especially when the robots are not in a uniform configuration. We propose a algorithm whereby each robots chooses another one to exchange only the measurements with so as to improve the tracking performance.

Simulation results validate the theoretical analysis showing that the self-triggered strategy converges to the optimal, uniform configuration. The average number of communication is less than 30% that of the constant strategy. In some cases, the self-triggered strategy converges faster than the constant communication strategy. We investigate possible causes. In particular, we observe that the robots travel a longer distance in the wrong direction in the constant strategy as opposed to the self-triggered strategy.

The rest of the paper is organized as follows. We start by formalizing the problem in Section 2. The self-triggered tracking strategy is presented in Section 3, assuming that the target's position is known and is fixed. We relax these assumptions and present a practical extension in Section 4. The simulation results are presented in Section 5. We conclude with a discussion of future work in Section 6.

A preliminary version of this paper was first presented in [20] without the observability based pairing strategy (Section 4.2) and the Gazebo simulation experiments (Section 5).

2 Problem Formulation

Consider a group of N robots moving on the boundary of a convex polygon $\mathcal{Q} \subset \mathbb{R}^2$. Let $\partial\mathcal{Q}$ denote the boundary of $\mathcal{Q}$. The robots are tasked with tracking a target o located in the interior of $\mathcal{Q}$. Let $p_1, ..., p_N$ denote the positions of the robots. We can map any point on $\partial\mathcal{Q}$ to a unit circle $\mathbb{T}$ using the transformation $\varphi_o : \partial\mathcal{Q} \to \mathbb{T}$ given by

$$\varphi_o(p) = \frac{p - o}{\|p - o\|} \tag{1}$$

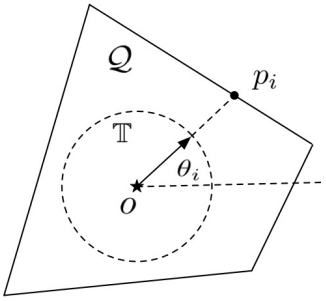

Fig. 1 The mapping from convex boundary $\partial\mathcal{Q}$ to unit circle $\mathbb{T}$.

as shown in Figure 1. We identify every robot's position with the corresponding point on the unit circle. That is, $p_i \in \partial\mathcal{Q} \subset \mathbb{R}^2$ is identified with $\theta_i = \varphi_o(p_i) \in \mathbb{T}$, indicating the location on the circle $\mathbb{T}$ of robot i. Let $\theta = (\theta_1, ..., \theta_N) \in \mathbb{T}^N$ denote the vector of locations of all robots.

We assume that all robots follow simple first-order continuous-time motion model. Let $\omega_{\max}$ denote the common maximum angular speed[1] for all robots on the unit circle. Our results can be extended to the situation where each robot has its own maximum angular speed.

Each robot i knows its own position exactly at all times. When two robots communicate they can exchange their respective positions.

Martinez and Bullo [13] showed that the optimal configuration for the robots that can obtain distance measurements towards the target is a uniform configuration along the circle where each pair of neighboring robots is equally spaced around the target. That is, $\theta_{i+1} - \theta_i = 2\pi/N, \forall i \in \{1, ..., N\}$. Optimality is defined with respect to maximizing the determinant of the Fisher Information Matrix (FIM). FIM is a commonly used measure for active target tracking. Martinez and Bullo [13] presented a decentralized control law that is guaranteed to (asymptotically) converge to a uniform configuration when a robot is allowed to communicate with only two of its immediate neighbors. That is, a robot i can communicate with only $i-1$ and $i+1$, along the circle. The analysis requires that all robots know the position of the target o exactly and that the target remains stationary. In the same paper, they showed how to apply the same control law in situations where the target's position is not known exactly and is instead estimated by combining noisy range measurements in an Extended Kalman Filter. They also evaluated the performance of the algorithm empirically in cases where the target is allowed to move.

[1] Strictly speaking, each robot has a maximum speed with which it can move on $\partial\mathcal{Q}$. In the appendix, we show how the maximum speed on $\partial\mathcal{Q}$ can be used to determine $\omega_{\max}$.

The control law in [13] assumes that neighboring robots communicate at every time step. We call this as the *constant strategy*. Our objective in this work is to reduce the number of communications between the robots while still maintaining the convergence properties. We present a *self-triggered strategy* where the control law for each robot not only decides how a robot should move, but also when should it communicate with its neighbors and seek new information. We show that the proposed self-triggered strategy is also guaranteed to converge to a uniform configuration, under the model and assumptions described in this section.

3 Self-Triggered Tracking Algorithm

In this section, we present the self-triggered tracking algorithm for achieving a uniform configuration along the unit circle. This requires knowing the center of the circle (*i.e.*, the target's position) and assuming that this center does not change. These assumptions are required for the convergence analysis to hold. We later relax these assumptions and present a practical version in the following section.

Our algorithm builds on the `self-triggered centroid algorithm` [15] which is a decentralized control law that achieves optimal deployment (*i.e.*, uniform Voronoi partitions) in a convex environment. We suitably modify this algorithm for the cases where the robots are restricted to move only on the boundary, $\partial\mathcal{Q}$, and can communicate with only two neighbors as described in the previous section. We first present the control law for each of the robot that uses the motion prediction set of its neighbors based on their last known positions. Then, we present an update policy to decide when a robot should communicate and seek new information from its neighbors.

3.1 Control Law

The constant control law in [13] drives every robot towards the midpoint of its Voronoi segment. The Voronoi segment of the robot i is the part of the unit circle extending from $(\theta_{i-1}+\theta_i)/2$ to $(\theta_i+\theta_{i+1})/2$. The constant control law steers robot i towards the midpoint of its Voronoi segment V^i_{mid} by using real-time (at every time-step) information from its neighbors, θ_{i-1} and θ_{i+1}, as illustrated in Figure 2. We refer to the book [16] for a comprehensive treatment on Voronoi segment.

In distributed self-triggered strategies, exact positions of the neighbors is not always available in real-time. Consequently, the algorithm must be able to operate on this inexact information. The information that each robot i holds about its neighbor j is the last known position of j, denoted by θ^i_j, and the time elapsed since the position of robot j was collected, denoted by τ^i_j. Based on this, robot i can compute the furthest distance that j could have moved in τ^i_j time as $\pm\phi^i_j$ where,

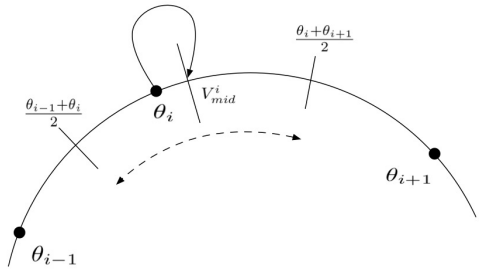

Fig. 2 Robot i goes towards the midpoint of its Voronoi segment via exact information from its neighbors.

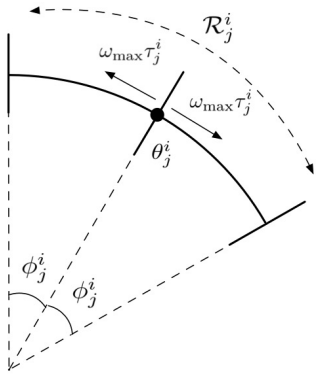

Fig. 3 Motion prediction set, $\mathcal{R}^i_j$, that each robot i maintains for its neighbors j. θ^i_j is the last known position of robot j and τ^i_j is the time elapsed since this last known position.

$$\phi^i_j = \omega_{\max}\tau^i_j. \tag{2}$$

Thus, robot i can build a prediction motion set $\mathcal{R}^i_j(\theta^i_j, \phi^i_j)$ that contains all the possible locations where robot j could have moved to in τ^i_j time (Figure 3).

Since robot i only communicates with its neighbors $i-1$ and $i+1$, the prediction motion range that robot i stores is given as
$\mathcal{R}^i := \{\mathcal{R}^i_{i-1}(\theta^i_{i-1}, \phi^i_{i-1}), \mathcal{R}^i_{i+1}(\theta^i_{i+1}, \phi^i_{i+1})\}$.

The proposed self-triggered strategy uses these motion prediction ranges $\mathcal{R}^i$ for defining the control law of robot i. Since the robot has inexact information of its neighbors, the midpoint of its Voronoi segment is a set instead of a point (Figure 4).

Define:

$$\theta^i_{i-1,\min} = \left(\theta^i_{i-1} - \phi^i_{i-1}\right) \quad \theta^i_{i-1,\max} = \left(\theta^i_{i-1} + \phi^i_{i-1}\right)$$
$$\theta^i_{i+1,\min} = \left(\theta^i_{i+1} - \phi^i_{i+1}\right) \quad \theta^i_{i+1,\max} = \left(\theta^i_{i+1} + \phi^i_{i+1}\right).$$

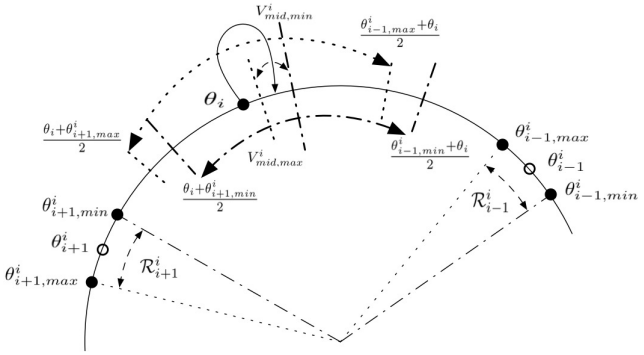

Fig. 4 Robot i goes towards the midpoint of its Voronoi segment via inexact motion prediction about its neighbors.

Thus, we have:

$$\mathcal{R}^i_{i-1}(\theta^i_{i-1}, \phi^i_{i-1}) = \{\beta \in \mathbb{T} | \theta^i_{i-1,\min} \leq \beta \leq \theta^i_{i-1,\max}\},$$
$$\mathcal{R}^i_{i+1}(\theta^i_{i+1}, \phi^i_{i+1}) = \{\beta \in \mathbb{T} | \theta^i_{i+1,\min} \leq \beta \leq \theta^i_{i+1,\max}\}.$$

Then the minimum and maximum midpoints of robot i's Voronoi segment can be computed as,

$$V^i_{\text{mid,min}} = \frac{(\theta^i_{i-1,\min} + \theta_i)/2 + (\theta_i + \theta^i_{i+1,\min})/2}{2}, \quad (3)$$

$$V^i_{\text{mid,max}} = \frac{(\theta^i_{i-1,\max} + \theta_i)/2 + (\theta_i + \theta^i_{i+1,\max})/2}{2}. \quad (4)$$

The midpoint of its Voronoi segment $V^i_{\text{mid}} \in [V^i_{\text{mid,min}}, V^i_{\text{mid,max}}]$. That is,

$$V^i_{\text{mid,min}} \leq V^i_{\text{mid}} \leq V^i_{\text{mid,max}}. \quad (5)$$

Substitute Equations 3 and 4 into Equation 5 yields,

$$\frac{\theta^i_{i+1} + 2\theta_i + \theta^i_{i-1} - 2\omega_{\max}\tau^i}{4} \leq V^i_{\text{mid}}$$

and

$$V^i_{\text{mid}} \leq \frac{\theta^i_{i+1} + 2\theta_i + \theta^i_{i-1} + 2\omega_{\max}\tau^i}{4},$$

then

$$-\frac{\omega_{\max}\tau^i}{2} \leq V^i_{\text{mid}} - \frac{\theta^i_{i+1} + 2\theta_i + \theta^i_{i-1}}{4} \leq \frac{\omega_{\max}\tau^i}{2}.$$

Therefore,

$$\left| V^i_{\text{mid}} - \frac{\theta^i_{i+1} + 2\theta_i + \theta^i_{i-1}}{4} \right| \leq \frac{\omega_{\max}\tau^i}{2}. \quad (6)$$

Thus, the angular distance between V^i_{mid} and $\frac{\theta^i_{i+1}+2\theta_i+\theta^i_{i-1}}{4}$ is bounded by $\frac{\omega_{\max}\tau^i}{2}$. In fact, the point $\frac{\theta^i_{i+1}+2\theta_i+\theta^i_{i-1}}{4}$ indicates the midpoint of i's guaranteed Voronoi segment gVs_i, defined as,

$$gVs_i = \left\{ \beta \in \mathbb{T} \left| \max_{\theta_i \in S_i} |\beta - \theta_i| \leq \min_{\theta_j \in S_j} |\beta - \theta_j|, \forall j \neq i \right. \right\}$$

where $T_1, \ldots, T_n \subset \mathbb{T}$ are a set of connected segments in $\mathbb{T}$. We refer to the report [5] for more details on the guaranteed Voronoi segment. Thus, the guaranteed Voronoi segment of robot i can be computed as,

$$gVs_i = \left\{ \beta \left| \frac{\theta_i + \theta^i_{i+1,\min}}{2} \leq \beta \leq \frac{\theta^i_{i-1,\max} + \theta_i}{2} \right. \right\}. \quad (7)$$

Although robot i does not know the exact midpoint of its Voronoi segment V^i_{mid}, it can move towards the midpoint of its guaranteed Voronoi segment gV^i_{mid} instead, which is given by,

$$gV^i_{\text{mid}} = \frac{(\theta_i + \theta^i_{i+1,\min})/2 + (\theta^i_{i-1,\max} + \theta_i)/2}{2},$$
$$= \frac{\theta^i_{i+1} + 2\theta_i + \theta^i_{i-1}}{4}. \quad (8)$$

In general, moving towards gV^i_{mid} does not guarantee that the robot moves closer to the midpoint of its Voronoi segment. However, the statement holds under the following condition.

Lemma 1 *Suppose robot i moves from θ_i towards gV^i_{mid}. Let θ'_i be its position after one time step. If $|\theta'_i - gV^i_{mid}| \geq |V^i_{mid} - gV^i_{mid}|$, then $|\theta'_i - V'_{mid}| \leq |\theta_i - V^i_{mid}|$.*

The proof for this lemma follows directly from the proof for Lemma 5.1 in [15]. Consequently, as long as the robot can ensure that its new position θ'_i satisfies $|\theta'_i - gV^i_{\text{mid}}| \geq |V^i_{\text{mid}} - gV^i_{\text{mid}}|$, then it is assured to not increase its distance from the actual (unknown) midpoint of the Voronoi segment. However, the right hand side of this condition also is not known exactly since robot i does not know V^i_{mid}. Instead, we can upper bound this term using Equation 6. We denote this upper bound by $\text{ubd}_i := \frac{\omega_{\max}\tau^i}{2}$. Thus, we get the following result:

Corollary 1 *Suppose robot i moves from θ_i towards gV^i_{mid}. Let θ'_i be its position after one time step. If*

$$|\theta'_i - gV^i_{mid}| > \text{ubd}_i, \quad (9)$$

then $|\theta'_i - V'_{mid}| \leq |\theta_i - V^i_{mid}|$.

Next, we present a motion control law that steers the robots towards a uniform configuration on the circle. Intuitively, robot i computes its guaranteed Voronoi segment (7) using the last known positions of its neighbors and the real-time position of itself. It then computes the midpoint of its guaranteed Voronoi segment (8) and moves towards the midpoint until it is within ubd_i of it. Formally, the control, $u_i(t_k)$, for robot i at time t_k is given by:

$$u_i(t_k) = \omega_i \text{ unit}(gV^i_{\text{mid}} - \theta_i), \quad (10)$$

where,

$$\omega_i = \begin{cases} \omega_{\max}, & \left|gV^i_{\text{mid}} - \theta_i\right| \geq \text{ubd}_i + \omega_{\max}\Delta t, \\ 0, & \left|gV^i_{\text{mid}} - \theta_i\right| \leq \text{ubd}_i, \\ \frac{\left|gV^i_{\text{mid}} - \theta_i\right| - \text{ubd}_i}{\Delta t}, & \text{otherwise.} \end{cases}$$

3.2 Triggering Policy

As time elapses, without new information the upper bound ubd_i grows larger until the condition in Equation 9 is not met. This triggers the robot to collect the updated information from its neighbors. There are two causes that may lead to the condition in Equation 9 being violated. The upper bound on the right hand side, ubd_i, might grow large because of the time elapsed since the last communication occurred. Or, robot i might move close to gV^i_{mid} which would require ubd_i to become small by acquiring new information. The second scenario might lead to frequent triggering when the robots are close to convergence. We introduce a user-defined tolerance parameter, $\sigma \geq 0$, to relax the triggering condition. Whenever the following condition is violated, the robot is required to trigger new communication:

$$\text{ubd}_i < \max\{\|\theta' - gV^i_{\text{mid}}\|, \sigma\} \tag{11}$$

Furthermore, the motion control law is designed under the assumption that the robot i and its two neighbors are located in the counterclockwise order. That is, $\theta_{i+1} > \theta_i > \theta_{i-1}$. Since the robots are identical, it is clear that there is no advantage gained by changing the order of robots along the circle. In a constant strategy, since the robots always communicate, they know the real-time position of their neighbors and can thus avoid the order being swapped. In a self-triggered strategy, however, we only have a motion prediction set of the neighbors. If there is a possibility that this order may be violated, the robots must communicate and avoid it. We achieve this by requiring the robot to maintain the following condition:

$$\theta^i_{i+1} - \omega_{\max}\tau^i_{i+1} > \theta_i > \theta^i_{i-1} + \omega_{\max}\tau^i_{i-1} \tag{12}$$

This ensures that even in the worst case, the robots have not swapped their positions. Whenever there is a possibility of this condition being violated, the robot i triggers a new communication.

The complete self-triggered midpoint strategy is presented below:

Algorithm 1: Self-triggered Midpoint

1: **while** all robots have not converged:
2: **for** each robot $i \in \{1, ..., N\}$ perform:
3: increment τ^i_{i-1} and τ^i_{i+1} by Δt
4: compute $\mathcal{R}^i, gVs_i, gV^i_{\text{mid}}$, and ubd_i
5: **if** Equation 11 OR Equation 12 is violated:
6: trigger communication with $i+1$ and $i-1$
7: reset τ^i_{i+1} and τ^i_{i-1} to zero
8: recompute $\mathcal{R}^i, gVs_i, gV^i_{\text{mid}}$, and ubd_i
9: **end if**
10: compute and apply u_i as defined in Equation 10
11: **end for**
12: **end while**

3.3 Convergence Analysis

Algorithm 1 is guaranteed to converge asymptotically to a uniform configuration along the circumference of the circle, irrespective of the initial configuration, assuming that no two robots are co-located initially. The proof for the convergence follows directly from the proof of Proposition 6.1 in [15] with suitable modifications. In the following, we sketch these modifications.

In [15] the robots are allowed to move anywhere in the interior of $Q \subset \mathbb{R}^2$ whereas in our case the robots are restricted to move on ∂Q, equivalent to moving on the unit circle $\mathbb{T}$. Therefore, all the L_2 distances in the proof in [15] change to L_1 distances. Instead of moving to the midpoint of the guaranteed Voronoi segment, the robots in [15] move to the centroid of a guaranteed Voronoi region. Instead of communicating with the two clockwise and counter-clockwise neighbors, the robots in [15] communicate with all possible Voronoi neighbors. None of these changes affect the correctness of the proof. We add an extra condition that triggers communication to prevent robots from changing their order along $\mathbb{T}$. Since this condition only results in additional triggers, it can only help convergence. Finally, since there is a one-to-one and onto mapping between ∂Q and $\mathbb{T}$, convergence along $\mathbb{T}$ implies convergence along ∂Q.

4 Practical Extension: Tracking of Moving Target with Noisy Measurements

If the true position of the target, o^*, is known, then we can draw a unit circle centered at the target and use the strategy in Algorithm 1 to converge to a uniform configuration along the circle. According to the result in [13] this configuration maximizes the determinant of the FIM. In practice, however, we do not know the true position of the target. In fact, the goal is to use the

noisy measurements from the robots to estimate the position of the target. Furthermore, the target may be mobile. This implies that the (unknown) center of the circle is also moving, further complicating the control strategy for the robots.

4.1 Decentralized Sensor Fusion

We use an Extended Kalman Filter (EKF) that estimates the position of the target (*i.e.*, center) and predicts its motion at every time step. The prediction and the estimate of the target from an EKF is a 2D Gaussian distribution parameterized by its mean, $\hat{o}_k$ and covariance. At each time step, we use the mean of the latest estimate as the center of the circle to compute the θ_i values using the transformation in Equation 1. We evaluate two regimes for the EKF:

- *Centralized EKF:* A common fusion center obtains the measurements from all the robots and computes a single target estimate, $\hat{o}_k$ at every time step; and
- *Decentralized EKF:* Each robot runs its own EKF estimator and has its own target estimate, $\hat{o}_k^i$, based on only its own measurements of the target. If at any time step, a robot communicates with its neighbors, then it also shares its current measurement with its neighbors. At these triggered instances, the robot updates its own estimate using its own measurement and measurements from its neighbors.

The rest of the process is similar to that in Algorithm 1. The centralized EKF scheme is a baseline which we compare against for the more realistic decentralized strategy. The results are presented in the simulation section that follows.

4.2 Pairing based on observability matrix

In the self-triggered strategy with decentralized EKF, between two successive triggered instances, robot i either uses its own measurements or those that of its immediate neighbors for estimating the target's state. If the neighbors are close to the robot then their measurements will be "similar" leading to poor estimate of the target's position (Figure 6(b). Specifically, since the robot obtains non-linear distance measurements, if robots $i+1, i, i-1$ are close to each other the resulting system gets closer to becoming unobservable.

In order to solve this problem, we propose a strategy that pairs each robot i with another robot j between two trigger instances (Figure 5). We assume that there is a central observer that knows the positions of each robot at all times and determines the optimal pairing

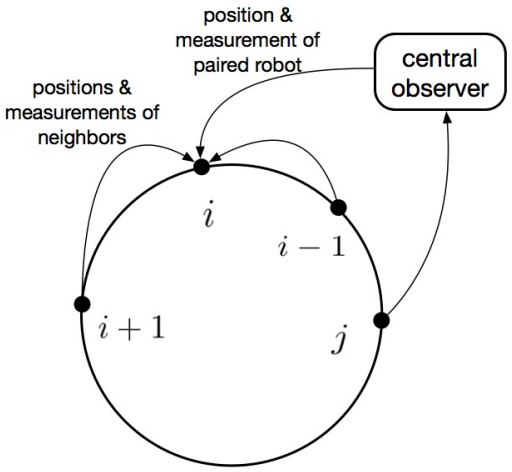

Fig. 5 Finding the optimal pair using the central observer. The neighbors $i-1$ and $i+1$ share their measurements and positions only at triggered time instances whereas the central observer shares the measurements and position of the paired robot j at all other times.

to improve the tracking performance. The central observer relays at all times to each robot, the position and measurement of its paired robot. The central observer does not need to relay any other information (including for example, the ordering of the paired robot or the number of robots). The criteria for the selection of robot j is based on the observability performance of the two-robot-target $(i-j-o)$ system:

$$\begin{cases} \dot{o} = u_o \\ z_i = \|p_i - o\|^2 \\ z_j = \|p_j - o\|^2 \end{cases} \tag{13}$$

where $o := [o_x, o_y]^T$ gives the 2D position of the target, and $u_o := [u_{ox}, u_{oy}]^T$ defines its control input, which is unknown and assumed to be fixed. z_i and z_j defines the range (or square of range)-only measurements from robot i and robot j of the target, with $p_i = [p_{ix}, p_{iy}]^T$ and $p_j = [p_{jx}, p_{jy}]^T$, respectively.

Using the Lie derivatives of sensor measurements, z_i, z_j, we construct the nonlinear observability matrix (see more details in [8]) for $i-j-o$ system as:

$$O^{i,j,o}(o, u_o) := \begin{bmatrix} O^{i,j,o}(o) \\ O^{i,j,o}(u_o) \end{bmatrix} = \begin{bmatrix} p_{ix} - o_x, p_{iy} - o_y \\ p_{jx} - o_x, p_{jy} - o_x \\ \\ O^{i,j,o}(u_o) \\ \\ \end{bmatrix} \tag{14}$$

where $O^{i,j,o}(o)$ and $O^{i,j,o}(u_o)$ defines the contribution from state and control input of the target, respectively. Here, since u_o is fixed, unknown and not controllable by

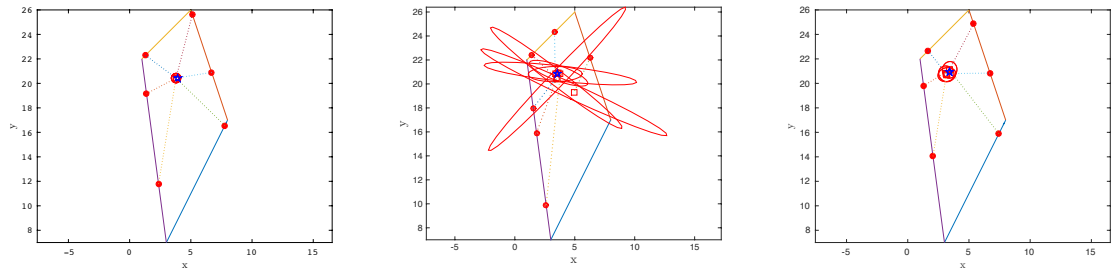

(a) (b) (c)

Fig. 6 (a) With a centralized EKF, all robots have the same estimate of the target leading to good convergence. (b) With measurements from self and immediate neighbors in decentralized EKF, the estimates of each robot are highly error-prone and uncertain. (c) This is improved by pairing each robot with another robot that improves the observability of the estimator.

the robot, we only focus on the state contribution part $O^{i,j,o}(o)$. Observability of trajectories (i.e., with control inputs) is not trivial (see for example [6] for one robot and one target case) and is the focus of our ongoing research.

The state of target o is *weakly locally observable* if the nonlinear observability matrix for $i-j-o$ system has full column rank [10]. However, the rank test for the observability of the system is a binary condition which does not tell the degree of the observability or how *good/bad* the observability is. We use the *condition number* of the observability matrix, defined as the ratio of the largest local singular value to the smallest, to measure this degree of unobservability [12]. A larger condition number suggests worse observability.

Based on this criterion, at the initial time instance, the central observer finds a pair j for each robot i, such that the condition number of the observability matrix of the $i-j-o$ system is minimized. Instead of recomputing the pairs at each time instance, we use a threshold to avoid repeated repairings. Once the pairs are set, the central observer will find a new pair only if the condition number for $i-j-o$ goes beyond a threshold. In the decentralized setting, the observer uses the estimate for robot i, i.e., $\hat{o}^i_k$ to compute the condition number when finding a pair for i at time k.

5 Simulation Results

In this section, we evaluate the performance of the proposed self-triggered tracking coordination algorithm. We first compare the convergence time for the self-triggered and constant communication strategies to achieve a uniform configuration on a convex boundary (Section 3). Then, we demonstrate the performance of the self-triggered and constant strategies for moving targets.

5.1 Stationary Target Case

In this section, we compare the performance of the self-triggered and constant strategies in terms of their convergence speeds and the number of communication messages to achieve a uniform configuration on the boundary of a convex environment. Here, we focus on the base case of known, stationary target position. All results are for 30 trials where the initial positions of the robots are drawn uniformly at random on the boundary.

Figure 7 shows snapshots of the active tracking process under the proposed self-triggered strategy starting with the initial configuration at time step $k=1$ in Figure 7-(a) and ending in a uniform configuration around the target at $k=760$ as shown in Figure 7-(c). For this example, we assume that the robots know the position of the stationary target. The initial positions of the robots are chosen uniformly at random on ∂Q. At each time step, we use the map φ_o to find θ_i on the unit circle (Equation 1), compute the control law as per Algorithm 1, and apply the inverse map φ^{-1} to compute the new positions of the robots on ∂Q. We set $\Delta t = 0.1\,s$ and assume that each robot has the same maximum angular velocity $\omega_{\max} = \frac{\pi}{180}\,\frac{rad}{s}$. In general, one can use the procedure given in the appendix to compute $\omega_{\max}$ for a given environment. Our MATLAB implementation is also available online.[2]

We first compare the convergence time of the two strategies (Figure 8-(a)). The convergence time, Ctime

[2] `https://github.com/raaslab/Self-triggered-mechanism`

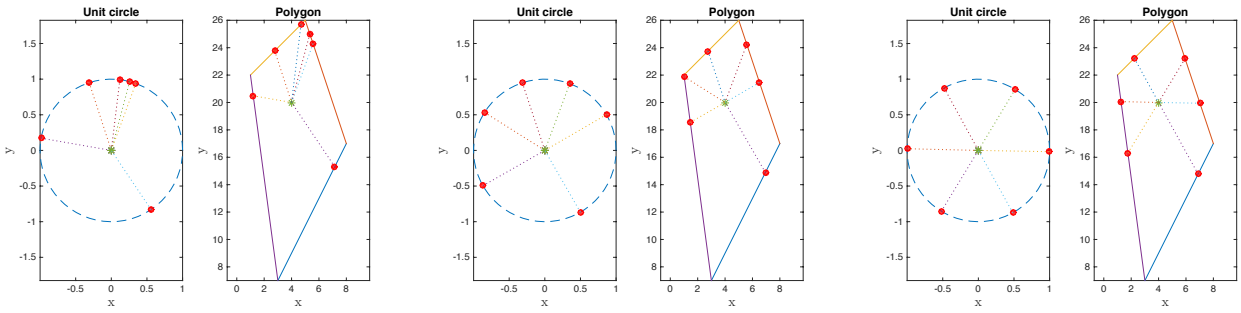

(a) $k = 1$ (b) $k = 400$ (c) $k = 760$

Fig. 7 Self-triggered tracking with six robots moving on the boundary of a convex polygon with a known, stationary target. The robots took 760 time steps to converge to the uniform configuration around the target.

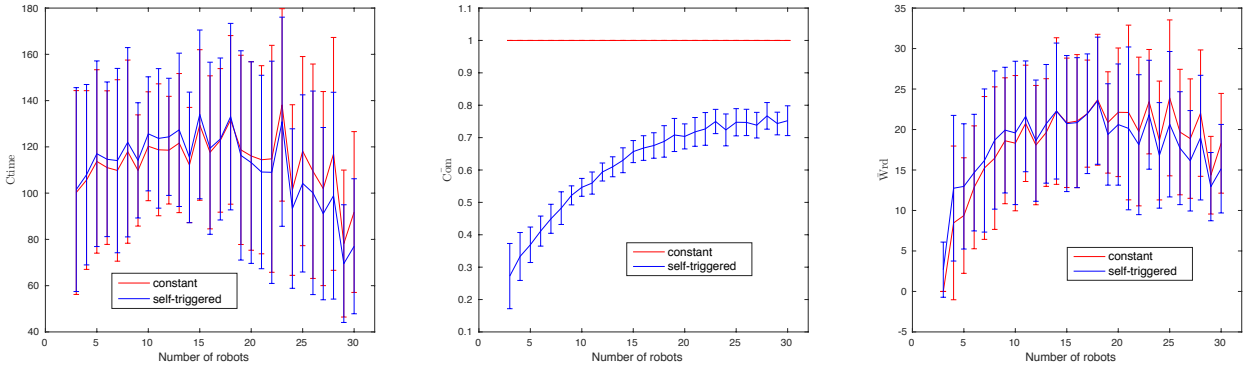

(a) (b) (c)

Fig. 8 Comparison of the convergence time (a) and the number of communication messages (b) and the average motion in the "wrong" direction (c) in constant and self-triggered strategies with a stationary target at known position.

is specified as the timestep k when the convergence error, Cerr, drops below a threshold. We use $0.1N$ as the threshold, where N is the number of robots. The convergence error term, Cerr, is defined as:

$$\text{Cerr} = \sum_{i=1}^{N} \left|\theta_i - V^i_{\text{mid}}\right| \tag{15}$$

in the constant communication case, and

$$\text{Cerr} = \sum_{i=1}^{N} \left|\theta_i - gV^i_{\text{mid}}\right| \tag{16}$$

in the self-triggered case.

The average number of communication messages is found as:

$$\overline{\text{Com}} = \frac{\sum_{i=1}^{N} \text{com}(i, \text{Ctime})}{N \times \text{Ctime}}$$

where $\text{com}(i, \text{Ctime})$ gives the total number of communications of a robot with its neighbors i at the end of Ctime. Figure 8-(b) shows the $\overline{\text{Com}}$ in the self-triggered case. The number of communication messages in the constant communication case is a constant.

Figure 8 shows that the self-triggered mechanism has faster convergence than the constant strategy when the number of robots is large. Intuitively, robots communicating constantly with its neighbors should converge faster. In order to investigate this counter-intuitive finding, we computed, $\overline{\text{Wrd}}$, which is the average motion in the "wrong" direction, defined as

$$\overline{\text{Wrd}} = \frac{\sum_{i=1}^{N} \text{wrd}(i, \text{Ctime})}{N \times \text{Ctime}}$$

Here, $\text{wrd}(i, \text{Ctime})$ gives the total amount of motion in the "wrong" directions for each robot i until the convergence time. For robot i, we use the sign of the difference between its initial orientation and the orientation at the convergence time, i.e., $\theta^0_i - \theta_i(\text{Ctime})$ to find which motion is in the "wrong" direction.

When the number of robots increases, the average distance traveled in the wrong direction becomes larger for the constant communication strategy. Recall that both strategies are guaranteed to converge only asymptotically. The rate of convergence is not known. We hypothesize that frequent communication with more robots, especially initially, leads to frequent switching in directions before the robots move towards their final configuration. On the other hand, in a self-triggered strategy the robots commit to a direction and move until the next triggered instance, thereby possibly leading to fewer switches in wrong direction.

We also implemented our algorithm in ROS and performed simulations in the Gazebo environment [11]. Figure 9 shows an instance with six differential-drive Pioneer 3DX robots [14] that can move in forwards and backwards direction.

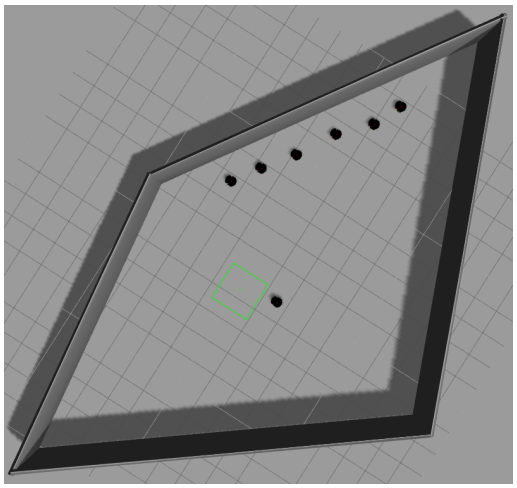

Fig. 9 Gazebo environment where six simulated Pioneer 3DX robots are tasked to track a target moving in the interior.

Figure 10-(a) shows that the constant communication strategy converges faster than the self-triggered one with six simulated robots. Changing the tolerance parameter σ affects the convergence time of the self-triggered strategy. The smaller the convergence tolerance σ, the faster the convergence which comes at the expense of increased number of messages. Figure 10-(b) shows communication messages for both strategies. Smaller the tolerance σ, larger the number of messages. The convergence tolerance σ acts as a trade-off between the communication messages and the convergence speed in self-triggered case.

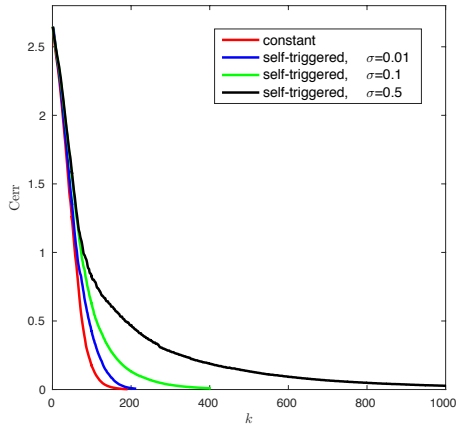

(a)

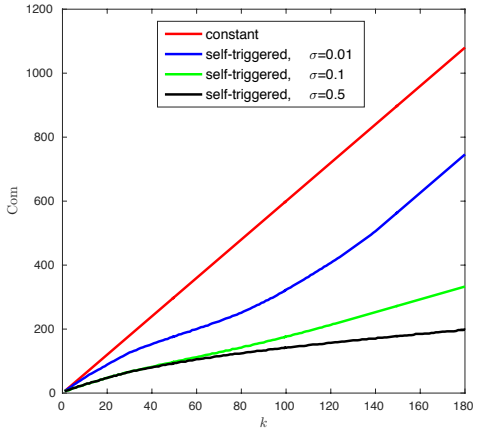

(b)

Fig. 10 Comparison of convergence error and communication messages in constant and self-triggered communication strategies using the setup shown in Figure 9.

5.2 Moving Target Case

Next, we present simulation results for the realistic case of mobile, uncertain targets (Section 4). We evaluate five strategies: constant communication with centralized EKF, constant communication with decentralized EKF, self-triggered communication with centralized EKF, self-triggered communication with decentralized EKF, and self-triggered communication with pairing decentralized EKF (Section 4.2). All five algorithms were implemented in Gazebo with six simulated Pioneer robots and a target moving on a circular trajectory.

Figure 11-(a),(b) show constant communication with centralized EKF and constant communication with decentralized EKF have similar tracking performance with respect to the convergence error, Cerr, over time. However, the target estimate error Terr is smaller in the cen-

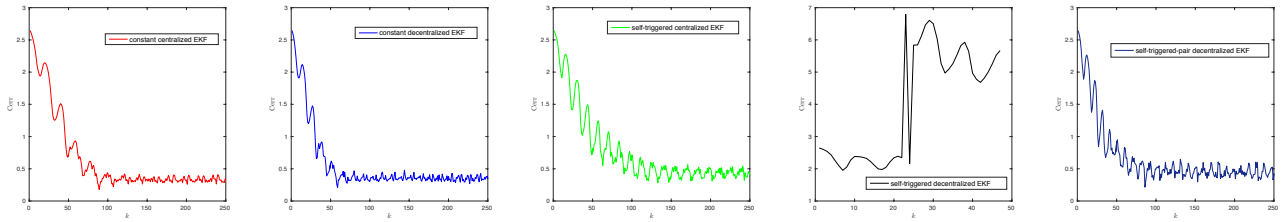

(a) (b) (c) (d) (e)

Fig. 11 Convergence error for mobile target tracking in constant communication with centralized EKF, constant communication with decentralized EKF, self-triggered communication with centralized EKF, self-triggered communication with decentralized EKF, and self-triggered communication with pairing decentralized EKF.

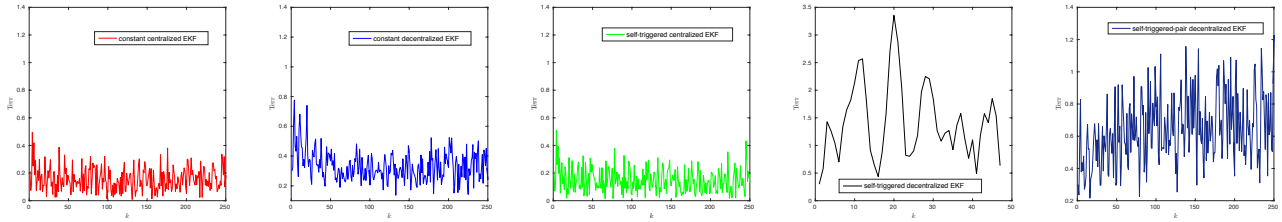

(a) (b) (c) (d) (e)

Fig. 12 Error in target's estimate for mobile target tracking in constant communication with centralized EKF, constant communication with decentralized EKF, self-triggered communication with centralized EKF, self-triggered communication with decentralized EKF, and self-triggered communication with pairing decentralized EKF.

tralized EKF case as shown in Figure 12-(a),(b). The target estimate error is defined as:

$$\text{Terr} = \|\hat{o} - o\|,$$

for the centralized case with $\hat{o}$ indicating the centralized estimate of the target, and

$$\text{Terr} = \frac{\sum_{i=1}^{N} \|\hat{o}_i - o\|}{N},$$

for the decentralized cases with $\hat{o}_i$ indicating the target estimate from each robot i. Figure 11-(d) and Figure 12-(d) show self-triggered strategy with decentralized EKF performs the worst in terms of both the convergence error and estimating the target's position. Using the paring mechanism based on the condition number of the observability matrix (Section 4.2) the tracking performance for the self-triggered strategy with decentralized EKF is improved (Figure 11-(e)).

Figure 13 shows the tracking performance of the self-triggered decentralized with pair strategy in relation to the baseline constant communication centralized EKF strategy as a function of the radius of the target's motion. We observe that the performance of the self-triggered strategy is comparable to the baseline algorithm.

6 Discussion and Conclusion

In this paper, we investigated the problem of active target tracking where each robot controls not only its own positions but also decide when to communicate and exchange information with its neighbors. We focused on a simpler target tracking scenario, first studied in reference [13]. We applied a self-triggered coordination strategy that asymptotically converges to a uniform configuration around the target while reducing the number of communication to less than 30% of a constant strategy. We find that the self-triggered strategy performs comparably with the constant communication strategy and in some cases even outperforms the baseline strategy. We conjecture that frequent communication in the constant communication strategy makes the robots subject to greedy and suboptimal than the less informative self-triggered strategy. Further investigating and proving this conjecture are part of our ongoing work. Future work also includes a extending the self-triggered strategy to decide not only when to communicate information, but also when to obtain measurements and which robots to communicate with. We conjecture that the latter question is crucial for better performance while tracking mobile targets.

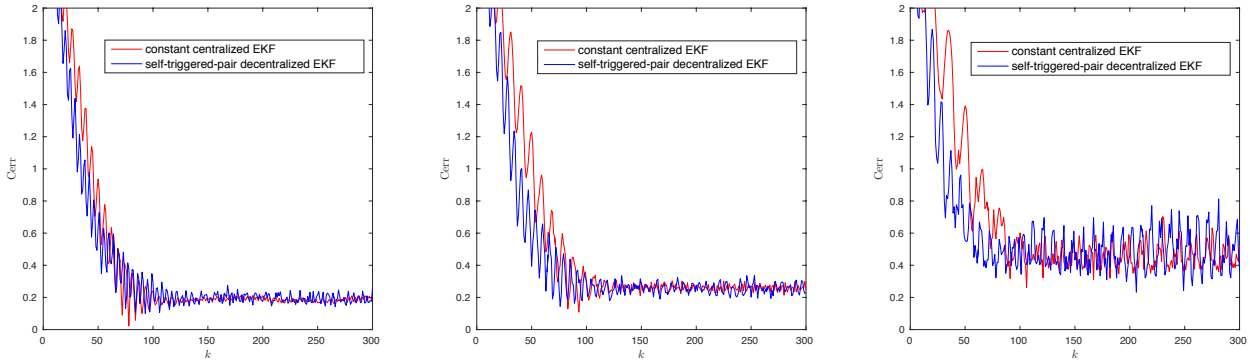

(a) $r = \frac{1.0m/s}{0.8rad/s}$ (b) $r = \frac{1.0m/s}{0.6rad/s}$ (c) $r = \frac{1.0m/s}{0.4rad/s}$

Fig. 13 Comparison of Cerr for mobile target tracking with constant centralized EKF and self-triggered-pair decentralized EKF w.r.t. the radius of the moving target.

References

1. Bar-Shalom, Y., Li, X.R., Kirubarajan, T.: Estimation with applications to tracking and navigation: theory algorithms and software. John Wiley & Sons (2004)
2. Dames, P., Tokekar, P., Kumar, V.: Detecting, localizing, and tracking an unknown number of moving targets using a team of mobile robots. In: International Symposium on Robotics Research (ISRR) (2015)
3. Dhillon, S., Chakrabarty, K.: Sensor placement for effective coverage and surveillance in distributed sensor networks, vol. 3. IEEE (2003)
4. Dunbabin, M., Marques, L.: Robots for environmental monitoring: Significant advancements and applications. IEEE Robotics and Automation Magazine **19**(1), 24 –39 (2012). DOI 10.1109/MRA.2011.2181683
5. Evans, W., Sember, J.: Guaranteed voronoi diagrams of uncertain sites (2008)
6. Gadre, A.S., Stilwell, D.J.: Toward underwater navigation based on range measurements from a single location. In: Robotics and Automation, 2004. Proceedings. ICRA'04. 2004 IEEE International Conference on, vol. 5, pp. 4472–4477. IEEE (2004)
7. Grocholsky, B., Keller, J., Kumar, V., Pappas, G.: Cooperative air and ground surveillance. Robotics & Automation Magazine, IEEE **13**(3), 16–25 (2006)
8. Hausman, K., Preiss, J., Sukhatme, G., Weiss, S.: Observability-aware trajectory optimization for self-calibration with application to uavs. arXiv preprint arXiv:1604.07905 (2016)
9. Heemels, W., Johansson, K.H., Tabuada, P.: An introduction to event-triggered and self-triggered control. In: 2012 IEEE 51st IEEE Conference on Decision and Control (CDC), pp. 3270–3285. IEEE (2012)
10. Hermann, R., Krener, A.: Nonlinear controllability and observability. IEEE Transactions on automatic control **22**(5), 728–740 (1977)
11. Koenig, N., Howard, A.: Design and use paradigms for gazebo, an open-source multi-robot simulator. In: Intelligent Robots and Systems, 2004.(IROS 2004). Proceedings. 2004 IEEE/RSJ International Conference on, vol. 3, pp. 2149–2154. IEEE (2004)
12. Krener, A.J., Ide, K.: Measures of unobservability. In: Decision and Control, 2009 held jointly with the 2009 28th Chinese Control Conference. CDC/CCC 2009. Proceedings of the 48th IEEE Conference on, pp. 6401–6406. IEEE (2009)
13. MartíNez, S., Bullo, F.: Optimal sensor placement and motion coordination for target tracking. Automatica **42**(4), 661–668 (2006)
14. Mei, Y., Lu, Y.H., Hu, Y.C., Lee, C.G.: A case study of mobile robot's energy consumption and conservation techniques. In: Advanced Robotics, 2005. ICAR'05. Proceedings., 12th International Conference on, pp. 492–497. IEEE (2005)
15. Nowzari, C., Cortés, J.: Self-triggered coordination of robotic networks for optimal deployment. Automatica **48**(6), 1077–1087 (2012)
16. Okabe, A., Boots, B., Sugihara, K., Chiu, S.N.: Spatial tessellations: concepts and applications of Voronoi diagrams, vol. 501. John Wiley & Sons (2009)
17. Rao, B., Durrant-Whyte, H.F., Sheen, J.: A fully decentralized multi-sensor system for tracking and surveillance. The International Journal of Robotics Research **12**(1), 20–44 (1993)
18. Tokekar, P., Branson, E., Vander Hook, J., Isler, V.: Tracking aquatic invaders: Autonomous robots for invasive fish. IEEE Robotics and Automation Magazine **20**(3), 33–41 (2013). DOI 10.1109/MRA.2012.2220506
19. Tokekar, P., Isler, V., Franchi, A.: Multi-target visual tracking with aerial robots. In: Proceedings of IEEE/RSJ International Conference on Intelligent Robots and Systems. IEEE (2014). DOI 10.1109/IROS.2014.6942986
20. Zhou, L., Tokekar, P.: Active target tracking with self-triggered communications. In: Proceedings of the IEEE International Conference on Robotics and Automation (ICRA) (2017). To appear

Calculation of $\omega_{\max}$

Assume robot has a maximum speed $v_{\max}$ with which it can move on $\partial\mathcal{Q}$. Thus, it can move as far as $d_{\max} = v_{\max}\Delta t$ in one time step Δt. We assume that $d_{\max}$ is less than the length of any edge of the polygon. Hence, a robot can cross at most one vertex per time step. Then we split the calculation of $\omega_{\max}$ into three separate cases (Figure 14).

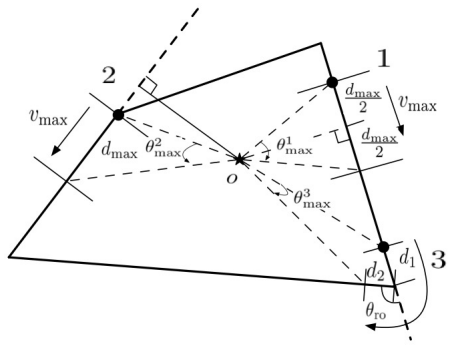

Fig. 14 Computing $\omega_{\max}$.

In all cases, let $\mathcal{E}^i$ be the edge on which the robot is located before moving a distance of $d_{\max}$. Let $l\mathcal{E}^i$ be the line supporting the edge. In cases 1 and 2, we compute $\omega_{\max}$ when the robot remains on $\mathcal{E}^i$ after traveling $d_{\max}$, whereas in case 3 the robot goes from $\mathcal{E}^i$ to $\mathcal{E}^{i+1}$

Case 1. The orthogonal projection of the target on $l\mathcal{E}^i$ lies within $\mathcal{E}^i$.

$\omega^{1,\mathcal{E}_i}_{\max}$ corresponds to the case where the robot covers a maximum angular distance with respect to the target in one time step. Thus, the robot should be as close as possible to the target when it moves $d_{\max}$ on the edge. $\omega^{1,\mathcal{E}_i}_{\max}$ can be calculated as $\omega^{1,\mathcal{E}_i}_{\max} = \frac{\theta^{1,\mathcal{E}_i}_{\max}}{\Delta t}$ giving $\omega^1_{\max} = \min_{\mathcal{E}_i \in \mathcal{E}}\{\omega^{1,\mathcal{E}_i}_{\max}\}$. Here $\theta^{1,\mathcal{E}_i}_{\max}$ is the angle shown in Figure 14.

Case 2. The orthogonal projection of the target on $l\mathcal{E}^i$ lies outside $\mathcal{E}^i$.

Similar to case 1, the $\omega_{\max}$ can be computed as $\omega^{2,\mathcal{E}_i}_{\max} = \frac{\theta^{2,\mathcal{E}_i}_{\max}}{\Delta t}$ where $\omega^2_{\max} = \min_{\mathcal{E}_i \in \mathcal{E}}\{\omega^{2,\mathcal{E}_i}_{\max}\}$. Here, $\theta^{2,\mathcal{E}_i}_{\max}$ is larger of the two angles made by the pair of lines joining target and either of the endpoint of $\mathcal{E}_i$ and joining target and a point $d_{\max}$ away from the corresponding endpoint.

Case 3. Robot crosses a vertex $\mathcal{V}_i$ within one time step.

We assume that within one time step Δt, robot moves $d_1^{\mathcal{V}_i}$ on one edge and $d_2^{\mathcal{V}_i}$ on another edge. Since robot must spend some time at the vertex turning in-place, we have $d_1^{\mathcal{V}_i} + d_2^{\mathcal{V}_i} < d_{\max}$. Thus, the $\omega^{3,\mathcal{V}_i}_{\max}$ can be calculated by the equation:

$$\frac{\theta^{3,\mathcal{V}_i}_{\max}}{\omega^{3,\mathcal{V}_i}_{\max}} = \frac{(d_1 + d_2)}{v_{\max}} + \frac{\theta^{\mathcal{V}_i}_{\text{ro}}}{\omega_{\text{ro}}} = \Delta t$$

where $\theta^{\mathcal{V}_i}_{\text{ro}}$ and ω_{ro} denote the rotation angle at the vertex $\mathcal{V}_i$ and rotational speed of the robot, respectively and $\theta^{3,\mathcal{V}_i}_{\max}$ is as shown in Figure 14. Then $\omega^3_{\max}$ can be specified as

$$\omega^3_{\max} = \min_{(\mathcal{V}_i, d_1)}\{\omega^{3,\mathcal{V}_i}_{\max}\}.$$

Where $\mathcal{V}_i \in \mathcal{V}$ and $0 \leq d_1 \leq \Delta t - \frac{\theta^{\mathcal{V}_i}_{\text{ro}}}{\omega_{\text{ro}}}$.

Finally, $\omega_{\max}$ can be computed as

$$\omega_{\max} = \min\{\omega^1_{\max}, \omega^2_{\max}, \omega^3_{\max}\}. \tag{17}$$

If $d_{\max}$ is larger than the length of one edge or the sum of lengths of several edges of the polygon, $\omega_{\max}$ can also be obtained using a similar procedure.